

Finding Mutated Subnetworks Associated with Survival in Cancer

Tommy Hansen*
tvhansen33@gmail.dk

Fabio Vandin^{†,*,‡,§}
vandinf@dei.unipd.it

Abstract

Next-generation sequencing technologies allow the measurement of somatic mutations in a large number of patients from the same cancer type. One of the main goals in the analysis of these mutations is the identification of mutations associated with clinical parameters, for example survival time. The identification of mutations associated with survival time is hindered by the genetic heterogeneity of mutations in cancer, due to the fact that genes and mutations act in the context of *pathways*. To identify mutations associated with survival time it is therefore crucial to study mutations in the context of interaction networks.

In this work we study the problem of identifying subnetworks of a large gene-gene interaction network that have mutations associated with survival. We formally define the associated computational problem by using a score for subnetworks based on the test statistic of the log-rank test, a widely used statistical test for comparing the survival of two given populations. We show that the computational problem is NP-hard and we propose a novel algorithm, called Network of Mutations Associated with Survival (NoMAS), to solve it. NoMAS is based on the color-coding technique, that has been previously used in other applications to find the highest scoring subnetwork with high probability when the subnetwork score is additive. In our case the score is not additive; nonetheless, we prove that under a reasonable model for mutations in cancer NoMAS does identify the optimal solution with high probability. We test NoMAS on simulated and cancer data, comparing it to approaches based on single gene tests and to various greedy approaches. We show that our method does indeed find the optimal solution and performs better than the other approaches. Moreover, on two cancer datasets our method identifies subnetworks with significant association to survival when none of the genes has significant association with survival when considered in isolation.

*Department of Mathematics and Computer Science, University of Southern Denmark, Odense (Denmark).

[†]Department of Information Engineering, University of Padova, Padova (Italy).

[‡]Department of Computer Science, Brown University, Providence, RI (USA).

[§]Corresponding author.

1 Introduction

Recent advances in next-generation sequencing technologies have enabled the collection of sequence information from many genomes and exomes, with many large human and cancer genetic studies measuring mutations in all genes for a large number of patients of a specific disease [8, 7, 5]. One of the main challenges in these studies is the interpretation of such mutations, in particular the identification of mutations that are clinically relevant. For example, in large cancer studies one is interested in finding somatic mutations that are associated with survival and that can be used for prognosis and therapeutic decisions. One of the main obstacles in finding mutations that are clinically relevant is the large number of mutations present in each cancer genome. Recent studies have shown that each cancer genome harbors hundreds or thousands of somatic mutations [16], with only a small number (e.g., ≤ 10) of *driver* mutations related to the disease, while the vast majority of mutations are *passenger*, random mutations that are accumulated during the process that leads to cancer but not related to the disease [41].

In recent years, several computational and statistical methods have been designed to identify driver mutations and distinguish them from passenger mutations using mutation data from large cancer studies [33]. Many of these methods analyze each gene in isolation, and use different single gene scores (e.g., mutation frequency, clustering of mutations, etc.) to identify significant genes [14, 24, 37]. While useful in finding driver genes, these methods suffer from the extensive *heterogeneity* of mutations in cancer, with different patients showing mutations in different cancer genes [22]. One of the reasons of such mutational heterogeneity is the fact that driver mutations do not target single genes but rather *pathways* [41], groups of interacting genes that perform different functions in the cell. Several methods have been recently proposed to identify significant groups of interacting genes in cancer [19, 25, 40, 26, 35, 23]. Many of these methods integrate mutations with interactions from genome-scale interaction networks, without restricting to already known pathways, that would hinder the ability to discover new important groups of genes.

In addition to mutation data, large cancer studies often collect also clinical data, including survival information, regarding the patients. An important feature of survival data is that it often contains *censored* measurements [21]: in many studies a patient may be alive at the end of the study or may leave the study before it ends, therefore only a lower bound to the survival of the patient is known. Survival information is crucial in identifying mutations that have a clinical impact. However, this information is commonly used only to assess the significance of candidate genes or groups of genes identified using other computational methods [19, 6], as the ones described above, and there is a lack of methods that integrate mutations, interaction information, and survival data to directly identify groups of interacting genes associated with survival.

The field of survival analysis has produced an extensive literature on the analysis of survival data, in particular for the comparison of the survival of two given populations (sets of samples) [21]. The most commonly used test for this purpose is the log-rank test [32, 28]. In genomic studies we are not given two populations, but a single set of samples, and are required to identify mutations that are associated with survival. The log-rank test can be used to this end to identify single genes associated with survival time by comparing the survival of the patients with a mutation in the gene with the survival of the patients with no mutation in the gene. The other commonly used test, the Cox Proportional-Hazards model [21], is equivalent to the log-rank test when the association of a binary feature with survival is tested, as it is in the case of interest to genomic studies. For a given group of genes, one can *assess* the association of mutations in the genes of the group with survival by comparing the survival of the patients having a mutation in at least one of the genes with the survival of the patients with no mutation in the genes. However, this approach cannot be used to *discover* sets of genes, since one would have to screen all possible subsets of genes and test their association with survival, and the number of subsets of genes to screen is enormous even considering only groups of genes interacting in a protein interaction network (e.g., there are $> 10^{15}$ groups of 8 interacting genes in HINT+HI2012 network [26]).

Color-coding is a probabilistic method that was originally described for finding simple paths, cycles and

other small subnetworks of size k within a given network [2]. The core of the color-coding technique is the assignment of random colors to the vertices, as a result of which the search space can be reduced, by restricting the subnetworks under consideration to *colorful* ones, those in which each vertex has a distinct color. For the identification of colorful subnetworks, dynamic programming is employed. The process is repeated until the desired subnetwork has been identified, that is having been colorful at least once, with high probability. When the dynamic programming algorithm is polynomial in n and the subnetworks being screened are of size $k \in O(\log n)$, the overall running time of the color-coding method too remains polynomial in n .

1.1 Our Contribution

In this paper we study the problem of finding sets of interacting genes with mutations associated to survival using data from large cancer sequencing studies and interaction information from a genome-scale interaction network. We focus on the widely used log-rank statistic as a measure of the association between mutations in a group of genes and survival. Our contribution is threefold: first, we formally define the problem of finding the set of k genes whose mutations show the maximum association to survival time by using the log-rank statistic as a score for a set of genes, and we show that such problem is NP-hard. We show that the problem remains hard when the set of k genes is required to form a connected subnetwork in a large graph with at least one node of large degree (*hub*).

Second, we propose an efficient algorithm, Network of Mutations Associated with Survival (NoMAS), based on the color-coding technique, to identify subnetworks associated with survival time. Color-coding has been previously used to find high scoring graphs for bioinformatics applications [12, 20] when the score for a subnetwork is *set additive* (i.e., the score of a subnetwork is the sum of the scores of the genes in the subnetwork). In our case the log-rank statistic is not set additive, and we prove that there is a family of instances for which our algorithm cannot identify the optimal solution. Nonetheless, we prove that under a reasonable model for mutations in cancer our algorithm identifies the optimal solution with high probability.

Third, we test our algorithm on simulated data and on data from three large cancer studies from The Cancer Genome Atlas (TCGA). On simulated data, we show that our algorithm does find the optimal solution while being much more efficient than the exhaustive algorithm that screens all sets of genes. On cancer data, we show that our algorithm finds the optimal solution for all values of k for which the use of the exhaustive algorithm is feasible, and identifies better solutions (in terms of association to survival) than a greedy algorithm similar to the one used in [34]. Moreover, we show that NoMAS identifies better solutions than using an (additive) score (i.e., the same gene score used in [38]) for a set of genes. For the cancer datasets, we show that our algorithm identifies novel groups of genes associated with survival where none of the genes is associated with survival when considered in isolation.

1.2 Related Work

Few methods have been proposed to identify groups of genes with mutations associated with survival in genomic studies. The work of [38] combines mutation and survival data with interaction information using a diffusion process on graphs starting from gene scores derived from p -values of individual genes, but did not consider the problem of directly identifying groups of genes associated with survival. The work of [34] combines mutation information and patient survival to identify subnetworks of a kinase-substrate interaction network associated with survival. It only focuses on phosphorylation-associated mutations, and the approach is based on a local search algorithm that builds a subnetwork by starting from one seed vertex and then greedily adding neighbours (at distance at most 2) from the seed, extending the approach used in different types of network analyses [11]. A similar greedy approach is used by [42] to identify groups of genes significantly associated with survival in cancer from gene expression data. For gene expression studies, [10] proposes an approach to enumerate dysregulated subnetworks in cancer based on an efficient search space pruning strategy, inspired by previous work on the identification of association rules in databases [36].

[31] uses the general approach described in [10] to identify subnetworks of genes with expression status associated to survival.

Color-coding has been previously used to count or search for subgraphs of large interaction networks ([1, 4]). Color-coding has also been used to identify groups of interacting genes in an interaction network that are associated with a phenotype of interest, but restricted to additive scores for sets of genes (i.e., the score of a set is the sum of the scores of the single genes); for example, [12] uses color-coding to find optimally discriminative subnetwork markers that predict response to chemotherapy from a large interaction network by defining a single gene score as $-\log_{10} d(g)$, where $d(g)$ is the discriminative score for gene g (i.e., a measure of the ability of g to discriminate two classes of patients); similarly, [20] uses color-coding to find groups of interacting genes with discriminative mutations in case-control studies, using as gene score the $-\log_{10}$ of the p -value from the binomial test of recurrence of mutations in the cases (while limiting the number of mutations in the controls).

2 Methods and Algorithms

In this section we define the model and the algorithm used in this work. The remaining of the section is organized as follows: Section 2.1 introduces preliminary definitions, the computational problem and presents its computational complexity; Section 2.2 introduces the algorithm we design to solve the problem; Section 2.3 presents the analysis of the algorithm, including the analysis under a reasonable model for mutations in cancer. Due to space constraints, proofs are omitted; they will appear in the final version of this extended abstract. (Proof sketches for our results are given in Appendix.)

2.1 Computational Problem

In survival analysis, we are given two populations (i.e., sets of samples) P_0 and P_1 , and for each sample $i \in P_0 \cup P_1$ we have its survival data: i) the survival time t_i and ii) the censoring information c_i , where $c_i = 1$ if t_i is the exact survival time for sample i (i.e., sample i is not censored), and $c_i = 0$ if t_i is a lower bound to the survival time for sample i (i.e., sample i is censored). Let m_0 be the number of samples in P_0 , m_1 be the number of samples in P_1 , and $m = m_0 + m_1$ be the total number of samples. Without loss of generality, the samples are $\{1, 2, \dots, m\}$, the survival times are $t = 1, 2, \dots, m$, with $t_i = i$ (i.e., the samples are sorted by increasing values of survival), and we assume that there are no ties in survival times. The survival data is represented by two vectors \mathbf{c} and \mathbf{x} , with c_i representing the censoring information for sample i , and x_i represents the population information: $x_i = 1$ if sample i is in population P_1 , and $x_i = 0$ otherwise. Given the survival data for two populations P_0 and P_1 , the significance in the difference of survival between P_0 and P_1 can be assessed by the widely used log-rank test [32, 28]. The log-rank statistic is

$$V(\mathbf{x}, \mathbf{c}) = \sum_{j=1}^m c_j \left(x_j - \frac{m_1 - \sum_{i=1}^{j-1} x_i}{m - j + 1} \right) \quad (1)$$

Under the (null) hypothesis of no difference in survival between P_0 and P_1 , the log-rank statistic asymptotically follows a normal distribution $\mathcal{N}(0, \sigma^2)$, where the standard deviation¹ is given by: $\sigma(\mathbf{x}, \mathbf{c}) =$

$$\sqrt{\frac{m_0 m_1}{m(m-1)}} \left(\left(\sum_{j=1}^m c_j \right) - \sum_{j=1}^m c_j \frac{1}{m-j+1} \right).$$

Thus the normalized log-rank statistic, defined as $\frac{V(\mathbf{x}, \mathbf{c})}{\sigma(\mathbf{x}, \mathbf{c})}$, asymptotically follows a standard normal $\mathcal{N}(0, 1)$ distribution, and the deviation of $\frac{V(\mathbf{x}, \mathbf{c})}{\sigma(\mathbf{x}, \mathbf{c})}$ from 0 is a measure of the difference in survival between P_0 and P_1 .

¹In the literature two different standard deviations (corresponding to two related but different null distributions, permutational and conditional) have been proposed for the normal approximation of the distribution of the log-rank statistic; we have previously shown [39] that the one we use here (corresponding to the permutational distribution) is more appropriate for genomic studies.

In genomic studies, we are given mutation data for a set \mathcal{G} of n genes in a set \mathcal{P} of m samples, represented by a mutation matrix M with $M_{i,j} = 1$ if gene i is mutated in patient j and $M_{i,j} = 0$ otherwise. We are also given survival data (survival time and censoring information) for all the m samples. Given a set $\mathcal{S} \subset \mathcal{G}$ of genes, we can assess the association of mutations in the set \mathcal{S} with survival by comparing the survival of the population $P_1^{\mathcal{S}}$ of samples with a mutation in at least one gene of \mathcal{S} and the survival of the population $P_0^{\mathcal{S}}$ of samples with no mutation in the genes of \mathcal{S} . That is, $P_0^{\mathcal{S}} = \{j \in \mathcal{P} : \sum_{i \in \mathcal{S}} M_{i,j} = 0\}$ and $P_1^{\mathcal{S}} = \{j \in \mathcal{P} : \sum_{i \in \mathcal{S}} M_{i,j} > 0\}$.

Given the set \mathcal{G} of all genes for which mutations have been measured, we are interested in finding the set $\mathcal{S} \subset \mathcal{G}$ with $|\mathcal{S}| = k$ that has maximum association with survival by finding the set \mathcal{S} that maximizes the absolute value of the normalized log-rank statistic. Given a set \mathcal{S} of genes, let $\mathbf{x}^{\mathcal{S}}$ be a 0 – 1 vector, with $x_i^{\mathcal{S}} = 1$ if at least one gene of \mathcal{S} is mutated in patient i , and $x_i^{\mathcal{S}} = 0$ otherwise. The normalized log-rank statistic for the set \mathcal{S} is then $\frac{V(\mathbf{x}^{\mathcal{S}}, \mathbf{c})}{\sigma(\mathbf{x}^{\mathcal{S}}, \mathbf{c})}$. Note that for a given set of patients the censoring information \mathbf{c} is fixed, therefore we can consider the log-rank statistic as a function $V(\mathbf{x}^{\mathcal{S}})$ of $\mathbf{x}^{\mathcal{S}}$ only. Analogously, we can rewrite $\sigma(\mathbf{x}^{\mathcal{S}}, \mathbf{c}) = \sigma(\mathbf{x}^{\mathcal{S}})f(\mathbf{c})$, where $\sigma(\mathbf{x}^{\mathcal{S}}) = \sqrt{m_1(m - m_1)}$ with $m_1 = |P_1^{\mathcal{S}}|$, and $f(\mathbf{c}) = \sqrt{\frac{1}{m(m-1)} \left(\left(\sum_{j=1}^m c_j \right) - \sum_{j=1}^m c_j \frac{1}{m-j+1} \right)}$ does not depend on $\mathbf{x}^{\mathcal{S}}$ and is fixed given \mathbf{c} .

To identify the set of k genes most associated with survival, we can then consider the score $|w(\mathcal{S})| = \left| \frac{V(\mathbf{x}^{\mathcal{S}})}{\sigma(\mathbf{x}^{\mathcal{S}})} \right|$. For ease of exposition in what follows we consider the score $w(\mathcal{S})$, corresponding to a one tail log-rank test for the identification of sets of genes with mutations associated with reduced survival; the identification of sets of genes with mutations associated with increased survival is done in an analogous way by maximizing the score $-w(\mathcal{S})$. We define the following problem.

The max k -set log-rank problem: Given a set \mathcal{G} of genes, an $n \times m$ mutation matrix M and the survival information (time and censoring) for the m patients in M , find the set $\mathcal{S} \subset \mathcal{G}$ of k genes maximizing $w(\mathcal{S})$.

We have the following.

Theorem 1. *The max k -set log-rank problem is NP-hard.*

We now define the max connected k -set log-rank problem that is analogous to the max k -set log-rank problem but requires feasible solutions to be connected subnetworks of a given graph \mathcal{I} , representing gene-gene interactions.

The max connected k -set log-rank problem: Given a set \mathcal{G} of genes, a graph $\mathcal{I} = (\mathcal{G}, E)$ with $E \subseteq \mathcal{G} \times \mathcal{G}$, an $n \times m$ mutation matrix M and the survival information (time and censoring) for the m patients in M , find the set \mathcal{S} of k genes maximizing $w(\mathcal{S})$ with the constraint that the subnetwork induced by \mathcal{S} in \mathcal{I} is connected.

If \mathcal{I} is the complete graph, the max connected k -set log-rank problem is the same as the max k -set log-rank problem. Thus, the max connected k -set log-rank problem is NP-hard for a general graph. However, we can prove that the problem is NP-hard for a much more general class of graphs.

Theorem 2. *The max connected k -set log-rank problem on graphs with at least one node of degree $O\left(n^{\frac{1}{c}}\right)$, where $c > 1$ is constant, is NP-hard.*

2.2 Algorithm

We design a new algorithm, Network of Mutations Associated with Survival (NoMAS), to solve the max connected k -set log-rank problem. The algorithm is based on an adaptation of the color-coding technique [2]. Our algorithm is analogous to other color-coding based algorithms that have been used before to identify subnetworks associated with phenotypes in other applications where the score is additive [12, 20].

The input to NoMAS is an undirected graph $G = (V, E)$, an $n \times m$ mutation matrix M , and the survival information \mathbf{x}, \mathbf{c} for the m patients in M . NoMAS first identifies a subnetwork \mathcal{S} with high score $\frac{w(\mathbf{x}^{\mathcal{S}})}{\sigma(\mathbf{x}^{\mathcal{S}})}$, and then uses a permutation test to assess the significance of the subnetwork.

To identify a subnetwork of high weight, the algorithm proceeds in iterations. In each iteration NoMAS colors G with k colors by assigning to each vertex v a color $\mathcal{C}(v) \in \{1, \dots, k\}$ chosen uniformly at random. For a given coloring of G , a subnetwork \mathcal{S} is said to be *colorful* if all vertices in \mathcal{S} have distinct colors. The *colorset* of \mathcal{S} is the set of colors of the vertices in \mathcal{S} . Note that the number of different colorsets (subsets of $\{1, \dots, k\}$) is 2^k . In each iteration the algorithm efficiently identifies high-scoring colorful subnetworks, and at the end the highest-scoring subnetwork among all iterations is reported.

Consider a given coloring of G . Let W be a $(2^k - 1) \times |V|$ table with a row for each non-empty colorset and a column for each vertex in G . Entry $W(T, u)$ stores the set of vertices of one connected colorful subnetwork that has colorset T and includes vertex u . Entries of W can be filled by dynamic programming. For colorsets of size 1, the corresponding rows in W are filled out trivially: $W(\{\alpha\}, u) = \{u\}$ if $\alpha = \mathcal{C}(u)$, and $W(\{\alpha\}, u) = \emptyset$ otherwise.

For entry $W(T, u)$ with $|T| \geq 2$, NoMAS computes $W(T, u)$ by combining a previously computed $W(Q, u)$ for u with another previously computed $W(R, v)$ where v is a neighbor of u in G , ensuring that the resulting subnetwork is connected and contains u . Colorfulness is ensured by selecting Q and R such that $Q \cap R = \emptyset$ and $Q \cup R = T$, and in turn ensures that $W(T, u)$ contains $|T|$ distinct vertices. Note that for a given T the choice of Q uniquely defines R . Thus, for each neighbor v of u there are (at most) $2^{|T|-1}$ possible combinations. Let $\mathcal{S}'(T, u)$ be the set of all colorful subnetworks that can be obtained by combining an entry $W(Q, u)$ for u and an appropriate entry $W(R, v)$ for a neighbor v of u so that $Q \cup R = T, Q \cap R = \emptyset$. That is:

$$\mathcal{S}'(T, u) = \bigcup_{\substack{v: (u,v) \in E \\ Q \cup R = T, Q \cap R = \emptyset}} \{W(Q, u) \cup W(R, v)\}.$$

(In the definition of $\mathcal{S}'(T, u)$ we assume that the union with \emptyset returns \emptyset .) $W(T, u)$ stores the element of $\mathcal{S}'(T, u)$ with largest value of our objective function, that is $W(T, u) = \arg \max_{\mathcal{S} \in \mathcal{S}'(T, u)} w(\mathcal{S})$. At the end, the best solution is identified by finding the entry of W of maximum weight. (See Appendix for pseudo code and illustrations of NoMAS).

After identifying the best solution \mathcal{S} for the mutation matrix M , NoMAS assesses its statistical significance by i) estimating the p-value $p(\mathcal{S})$ for the log-rank statistic (using a Monte-Carlo estimate with 10^8 samples), and then ii) using a permutation test in which \mathcal{S} is compared to the best solution \mathcal{S}^p for the mutation matrix M^p obtained by randomly permuting the rows of M . A total of 100 permutations are performed and the *permutation* p-value is recorded as the ratio of permutations in which $w(\mathcal{S}^p) \geq w(\mathcal{S})$. While the p-value from the log-rank test reflects the association between mutations in the subnetwork and survival, the permutation p-value assesses whether a subnetwork with association with survival at least as extreme as the one observed in the input data can be observed when the genes are placed randomly in the network. Note that we can identify multiple solutions by considering different entries of W (even if the same solution may appear in multiple entries of W), and we obtain a permutation p-value for the i -th top scoring solution by comparing its score with the score of the i -th top scoring solution in the permuted datasets. Analogously, NoMAS identifies sets that minimize $w(\mathcal{S})$ (sets associated to increased survival) by maximizing the score $-w(\mathcal{S})$.

Parallelization. The computation of W is parallelized using $N \leq |V|$ processors. All entries of W are kept in shared memory and $|V|/N$ unique columns uniformly at random are assigned to each processor.

Entries of W are computed in order of increasing colorset sizes. We define the i -th *colorset group* as the set of all $\binom{k}{i}$ colorsets of size i . We exploit the fact that rows within the i -th colorset group are computed by reading entries exclusively from rows belonging to colorset groups $< i$. When a processor has finished the rows of the i -th colorset group it waits for the other processors to do the same. When the last processor

completes the i -th colorset group, all N processors can safely begin to compute rows of colorset group $i + 1$. In total, k synchronization steps are carried out, one for each colorset group.

2.3 Analysis of NoMAS

We consider the performance of NoMAS excluding the permutation test. The log-rank statistic $w(\mathcal{S})$ is computed in time $O(m_1) \in O(m)$. The total time complexity for computing a single entry $W(T, u)$ is then bounded by $O(m \deg(u) 2^{|T|-1}) \in O(m \deg(u) 2^k)$, where $\deg(u)$ is the degree of u in G . Given a coloring of G , the computation of the entire table can thus be performed in time $O(2^k \sum_{u \in V} m \deg(u) 2^k) \in O(m|E|4^k)$. If L iterations are performed, then the complexity of the algorithm is $O(Lm|E|4^k)$.

Let OPT be the optimal solution. If the score $w(\mathcal{S})$ was set additive, as the scores considered in previous applications of color-coding for optimization problems on graphs, to discover OPT it would be sufficient that OPT be colorful, that happens with probability $k!/k^k \geq e^{-k}$ for each random coloring. Therefore $O(\ln(1/\delta)e^k)$ iterations would be enough to ensure that the probability of OPT not being discovered is $\leq \delta$, resulting in an overall time complexity of $O(m \ln(1/\delta)|E|(4e)^k)$.

However, our score $w(\mathcal{S})$ is not set additive (e.g., if two genes in \mathcal{S} have a mutation in the same patient the weight of the patient is considered only once in $w(\mathcal{S})$). Therefore, while OPT being colorful is still a necessary condition for the algorithm to identify OPT , the colorfulness of OPT is not a sufficient condition. In fact, we have the following.

Proposition 1. *For every $k \geq 3$ there is a family of instances of the max connected k -set log-rank problem and colorings for which OPT is not found by NoMAS when it is colorful.*

Even more, we prove that when mutations are placed arbitrarily then for every subnetwork \mathcal{S} and a given coloring of \mathcal{S} , any color-coding algorithm that adds subnetworks of size k to W by merging neighboring subnetworks of size $< k$ could be “fooled” to not add \mathcal{S} to W by simply adding 3 vertices to G and assigning them a specific color.

Theorem 3. *For any optimal colorful connected subnetwork \mathcal{S} of size $k \geq 3$ and any color-coding algorithm \mathcal{A} which obtains subnetworks with colorsets of cardinality i by combining 2 subnetworks with colorsets of cardinality $< i$, by adding 3 neighbors to \mathcal{S} we have that \mathcal{A} may not discover \mathcal{S} .*

Intuitively, Proposition 1 and Theorem 3 show that if mutations are placed adversarially (and the optimal solution OPT has many neighbors), our algorithm may not identify OPT . However, we prove that our algorithm identifies the optimal solution under a generative model for mutations, that we deem the *Planted Subnetwork Model*. We consider $w(\mathcal{S})$ as the unnormalized version of the log-rank statistic. In this model: i) there is a subnetwork \mathcal{D} , $|\mathcal{D}| = k$, with $w(\mathcal{D}) \geq cm$, for a constant $c > 0$; ii) each gene $g \in \mathcal{D}$ is such that $w(\mathcal{D}) - w(\mathcal{D} \setminus \{g\}) \geq \frac{c'm}{k}$, for a constant $c' > 0$; iii) for each gene $g \in \mathcal{D}$: $w(\{g\}) > 0$; iv) for each gene $\hat{g} \notin \mathcal{D}$, \hat{g} is mutated with probability p_g in each patient, independently of all other events (and of survival time and censoring status).

Intuitively, i) above states that the subnetwork \mathcal{D} has mutations associated with survival; ii) states that each gene $g \in \mathcal{D}$ contributes to the association of mutations in \mathcal{D} to survival; iii) states that each gene $g \in \mathcal{D}$ should have the same association to survival (increased or decreased) as \mathcal{D} ; and iv) states that all mutations outside \mathcal{D} are independent of all other events (including survival time and censoring of patients).

We show that when enough samples are generated from the model above, our algorithm identifies the optimal solution with the same probability guarantee given by the color-coding technique for additive scores.

Theorem 4. *Let M be a mutation matrix corresponding to m samples from the Planted Subnetwork Model. If $m \in \Omega(k^4(k + \varepsilon) \ln n)$ for a given constant $\varepsilon > 0$ and $O(\ln(1/\delta)e^k)$ color-coding iterations are performed, then our algorithm identifies the optimal solution \mathcal{D} to the max connected k -set log-rank with probability $\geq 1 - \frac{1}{n^\varepsilon} - \delta$.*

3 Results

We assessed the performance of NoMAS by using simulated and cancer data. We compared NoMAS to the exhaustive algorithm that identifies the subnetwork of k vertices with the highest score $w(S)$ for the values of k for which we could run the exhaustive algorithm (we implemented a parallelized version of the algorithm described in [29] to efficiently enumerate all connected subnetworks), to three variants of a greedy algorithm similar to the one from [34], and to the use of a score given by the sum of single gene scores. Cancer data is obtained from The Cancer Genome Atlas (TCGA). In particular, we consider somatic mutations (single nucleotide variants and small indels) for 268 samples of glioblastoma multiforme (GBM), 315 samples of ovarian adenocarcinoma (OV), and 174 samples of lung squamous cell carcinoma (LUSC) for which survival data is available.

For all our experiments we used as interaction graph G the graph derived from the application of a diffusion process on the HINT+HI2012 network², a combination of the HINT network [13] and the HI-2012 [44] set of protein-protein interactions, previously used in [25]. The details of the diffusion process are described in [25]. In brief, for two genes g_i, g_j the diffusion process gives the amount of heat $h(g_i, g_j)$ observed on g_j when g_i has one mutation, and the amount of heat $h(g_j, g_i)$ observed on g_i when g_j has one mutation. The graph used for our analyses is obtained retaining an edge between g_i and g_j if $\max\{h(g_i, g_j), h(g_j, g_i)\} \geq 0.012$. The resulting graph has 9859 vertices and 42480 edges, with the maximum degree of a node being 438. In all our experiments we removed mutations in genes mutated in < 3 of the samples. For cancer data, this resulted in 780 mutated genes in GBM, 890 in OV, and 2915 in LUSC.

The machine, on which all our experiments were carried out, consists of two CPUs of the type Intel Xeon E5-2698 v3 (2.30GHz), each with 16 physical cores, for a total of 64 virtual cores, and 16 banks of Samsung 32GB DDR4 (2133 MHz) memory modules for a total of 512GB of memory.

The remaining of the section is organized as follow: Section 3.1 presents the results on simulated data, while Section 3.2 presents the results on cancer data.

3.1 Simulated data

We assess the performance of NoMAS on simulated data generated under the Planted subnetwork Model. The subnetwork $\mathcal{D} \subset \mathcal{G}, |\mathcal{D}| = k$ associated with survival is generated by a random walk on the graph G . We model the association of \mathcal{D} to survival by mutating with probability p one gene of \mathcal{D} chosen uniformly at random in each sample among the $\frac{m}{4}$ of lowest survival. All other genes in \mathcal{D} are mutated independently with probability 0.01 in all samples, to simulate passenger mutations (not associated with survival) in \mathcal{D} [24]. For genes in $\mathcal{G} \setminus \mathcal{D}$, we used the same mutation frequencies observed in the GBM study, and mutate each gene independently of all other events.

We fixed $k = 5$ and considered the values of $p \in \{0.5, 0.75, 0.85\}$ and $m \in \{268, 500, 750, 1000\}$. We kept the same ratio of censored observations as in GBM and chose the censored samples uniformly among all samples. For every pair (p, m) we performed 100 simulations, running NoMAS on the dataset with $L = 256$ color-coding iterations, and recorded whether NoMAS reported \mathcal{D} as the highest scoring subnetwork. Results are shown in Fig. 1 (a). For sample sizes similar to the currently available ones, NoMAS frequently reports \mathcal{D} as the highest scoring solutions when there is a quite strong association of \mathcal{D} with survival ($p \geq 0.85$), but for $m = 1000$ the highest scoring subnetwork reported by NoMAS is \mathcal{D} in $> 80\%$ of the cases even for $p = 0.5$. Fig. 1(b) shows that even when NoMAS does not report \mathcal{D} as the highest scoring solution, the solution reported by NoMAS contains mostly genes that are in \mathcal{D} , even for current sample size (e.g., on average 74% of the genes in the \mathcal{D} are reported by NoMAS for $m = 268$ and $p = 0.85$ even when \mathcal{D} is not the highest scoring solution by NoMAS). Finally, we assessed whether \mathcal{D} would be among the highest scoring solutions in the table W computed by NoMAS: Fig 1(c) shows that by considering the top-10 solutions W the chances to identify \mathcal{D} increase substantially even for $m = 268$

²<http://compbio-research.cs.brown.edu/pancancer/hotnet2/>

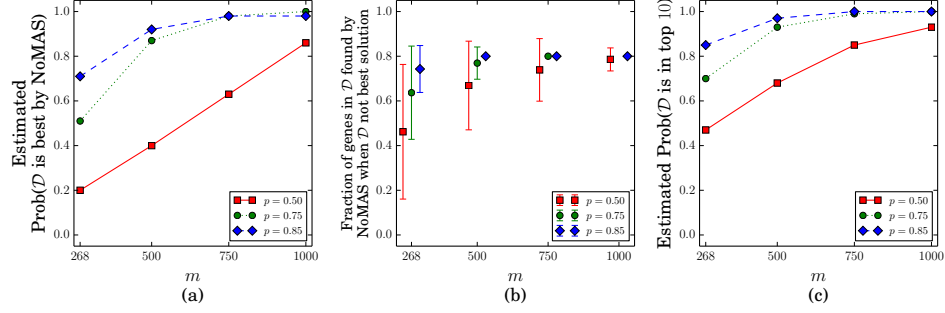


Figure 1: Results of NoMAS on simulated data from the Planted Subnetwork Model. 100 datasets were generated for each pair (m, p) , where m is the number of samples and for different probabilities p of mutations in the set \mathcal{D} of genes associated with survival. (a) Probability that \mathcal{D} is reported as the highest scoring solution by NoMAS. (b) Ratio of genes from the set \mathcal{D} that are in the best solution when \mathcal{D} is not the highest scoring solution by NoMAS. (c) Probability that \mathcal{D} is among the top-10 solutions reported by NoMAS. All probabilities are estimated from the simulated datasets.

and $p = 0.5$, with most configurations having > 0.8 probability of finding \mathcal{D} in the top-10 solutions by NoMAS. For a fixed $p = 0.75$ and for each value of m we assessed whether NoMAS identified the optimal solution even when it was not \mathcal{D} (an event not excluded in the Planted subnetwork Model) and found that for $m \geq 500$ NoMAS reported the optimal solution in 10 out of 10 cases (for $m = 268$ NoMAS identified the optimal solution 9 out of 10 times). These results show that NoMAS does indeed find the optimal solution in almost all cases even for sample sizes currently available (while the theoretical analysis of Section 2.3 suggests that much larger sample sizes are required) and it can be used to identify \mathcal{D} or the majority of it by considering the top-10 highest scoring solutions.

The performance of NoMAS is affected when altering the ratio of samples that, with probability p , are mutated in \mathcal{D} . A higher ratio results in an increased performance of NoMAS, i.e. more cases arise in which \mathcal{D} is the best subnetwork identified, and equally is \mathcal{D} more frequently in the top 10. When increasing the ratio (for example $\frac{m}{3}$ in stead of $\frac{m}{4}$) more samples are being mutated, and each gene in \mathcal{D} receives more mutations (the gene responsible for each of the mutated samples is chosen uniformly at random). The increment in the association to survival is thus increased each time one of the genes is added to \mathcal{D} .

3.2 Cancer data

We assessed the performance of NoMAS on the GBM, OV, and LUSC datasets. We first assessed whether NoMAS identified the optimal solution by comparing the highest scoring solution reported by NoMAS with the one identified by using the exhaustive algorithm for $k = 2, 3, 4, 5$. In all cases we found that NoMAS does identify the optimal solution, while requiring much less running time compared to the exhaustive algorithm (Supplementary Fig. 2). For $k > 5$ we could not run the exhaustive algorithm, while the runtime of NoMAS is still reasonable. The runtime of NoMAS can be greatly improved by using the parallelization strategy described in Section 2.2 (Supplementary Fig. 3). We therefore used NoMAS to find subnetworks of size $k = 6$ and $k = 8$. We also considered two modifications of NoMAS that solve some easy cases where NoMAS may not identify the highest scoring solution due to its subnetwork merging strategy (see Appendix for a description and pseudo code of the modifications). We run both modifications on GBM, OV, and LUSC for $k = 6, 8$ (using the same colorings used by the original version of NoMAS): in all cases the modified versions of NoMAS did not report subnetworks with higher scores than the ones from the original version of NoMAS. We also note that the original version of NoMAS is significantly faster in practice than its two modifications (Supplementary Fig 3), and therefore we used the original version of NoMAS in the remaining experiments.

We also compared NoMAS with three different greedy strategies for the max connected k -set log-rank

problem. All three algorithms build solutions starting from each node $u \in G$ in iterations by adding nodes to the current solution \mathcal{S} , and differ in the way they enlarge the current subnetwork \mathcal{S} of size $1 \leq i < k$. (See Appendix for a description of the greedy algorithms). We run the three greedy algorithms on GBM, OV, and LUSC for $k = 4, 5, 6, 8$. For each dataset we compared the resulting subnetworks with the ones identified by NoMAS. Results are shown in Fig. 2. In almost all cases we found that NoMAS discovered subnetworks with higher score than the subnetworks found by using greedy strategies, even if in some cases there is a greedy strategy that identifies the same subnetworks for all values of k . The difference in score increases as k increases, showing the ability of NoMAS to discover better solutions for larger values of k (see Supplementary Fig. 4 for a running time comparison between NoMAS and the greedy algorithms). We also assessed whether the fact that greedy strategies discover lower scoring solutions than NoMAS has an impact on the estimate of the p -value in the permutational test. We considered the top-10 scoring solutions (corresponding to 10 different starting nodes $u \in G$) discovered by the best greedy strategy in the GBM dataset, and compute the permutational p -value for each solution by generating 100 permuted datasets and either use the (same) greedy strategy for permuted data or use NoMAS for permuted data (using only 32 iterations on the permuted data) Supplementary Fig. 1 shows a comparison of the distribution of the p -values. As we can see, the greedy strategy incorrectly underestimate the permutational p -values for the solutions, due to the greedy algorithm not being able to identify solutions of score as high as NoMAS in the permuted datasets. The use of the greedy algorithms would then lead to both 1. identify solutions in real data with lower association to survival compared to NoMAS and 2. wrongly estimate their permutational p -value as more significant than it is.

Finally, we compared NoMAS with the use of an (additive) score that sums single gene scores (similar to the ones used in [38]). For each gene $g \in G$ we computed the p -value $p(g)$ for the association of g with survival using the log-rank test, and define $a(\mathcal{S}) = \sum_{g \in \mathcal{S}} -\log_{10} p(g)$. We then partitioned the genes according to their association with increased survival or with decreased survival, and modified our algorithm to look for high scoring solutions in a partition using score $a(\mathcal{S})$. Results are in Fig 2. We found that NoMAS outperforms the use of a single gene score, with a very large difference for certain values of the parameters.

We then considered the top-10 highest scoring subnetworks obtained from NoMAS on GBM and OV for $k = 8$. For each subnetwork \mathcal{S} we estimated the log-rank p -value and the permutational p -value as described in Section 2.2. These subnetworks do not contain any *gene* that would be reported as significant by single gene tests (corrected $p=1$), but they all show a high association with survival: in GBM, all subnetwork have log-rank p -value $< 2 \times 10^{-6}$; in OV all subnetwork have p -value $< 5 \times 10^{-7}$.

For example, in OV we identify a subnetwork (Fig 3a) of 8 genes associated with increased survival including BRCA2, a known cancer gene previously reported to have mutations associated with improved survival [43] in ovarian cancer, MSTR1, previously reported as a novel prognostic marker in gastroesophageal adenocarcinoma [9], and MYH9, associated to metastasis and tumor invasion in gastric cancer [27]. The permutational p -value for this subnetwork is 0.07, and therefore it is unlikely this association is due to random variation. In GBM, we identify a subnetwork (Fig 3b; permutational p -value = 0.2) of 8 genes associated to decreased survival, including CDKN2A, an important component of the Rb pathway in GBM [3], BID, a known cell death regulator, SCAP, involved in the molecular mechanisms of lipid metabolism in gliomas [17], TRPV4, previously identified as part of the cell migration mechanism [15], CARD6, involved in the antiapoptotic pathway and previously found to be associated to drug sensitivity in human glioblastoma multiforme cell lines [18], and IRS1, part of the known PI3K/PTEN cancer pathway [30]. These results show that NoMAS identifies subnetworks associated with survival data in GBM and OV including known cancer genes and genes previously reported to be associated with survival, as well as genes that represent novel candidates for the association with survival.




Figure 2: Comparison of the normalized log-rank statistic of the best solution reported by NoMAS, by greedy algorithms (see Appendix for the description), and by the algorithm that uses an additive scoring function $a(\mathcal{S})$ (denote by “additive” in the plots). To maintain readability we omit values above -4.0 when considering mutations associated with increased survival. For each datasets the results for the maximization of $w(\mathcal{S})$ (top panel) and the maximization of $-w(\mathcal{S})$ (bottom panel) are shown separately. (a) Results for GBM dataset. (b) Results for OV dataset. (c) Results for LUSC dataset.

4 Conclusion

In this work we study the problem of identifying subnetworks of a large gene-gene interaction network that are associated with survival using mutations from large cancer genomic studies. We formally define the associated computational problem, that we call the max connected k -set log-rank problem, by using as score for a subnetwork the test statistic of the log-rank test, one of the most widely used statistical tests to assess the significance in the difference in survival among two populations. We prove that the max connected k -set log-rank problem is NP-hard in general, and is NP-hard even when restricted to graphs with at least one node of large degree. We develop a new algorithm, NoMAS, based on the color-coding technique, to efficiently identify high-scoring subnetworks associated with survival. We prove that even if our algorithm is not guaranteed to identify the optimal solution with the probability given by the color-coding technique (due the non additivity of our scoring function), it does identify the optimal solution with the same guarantees given by the color-coding technique when the data comes from a reasonable model for mutations and independently of the survival data. Using simulated data, we show that NoMAS is more efficient than the exhaustive algorithm while still identifying the optimal solution, and that our algorithm will identify subnetworks associated with survival when sample sizes larger than most currently available ones, but still reasonable, are available.

We use cancer data from three cancer studies from TCGA to compare NoMAS to approaches based on single gene scores and to greedy methods similar to ones proposed in the literature for the identification of subnetworks associated with survival and for other problems on graphs. Our results show that NoMAS identify subnetworks with stronger association to survival compared to other approaches, and allows the correct estimation of p -values using a permutation test. Moreover, in two datasets NoMAS identifies two subnetworks associated with survival containing genes previously reported to be important for prognosis in




Figure 3: Two subnetworks identified by NoMAS on real cancer data. (a) Subnetwork S associated with survival in GBM, and Kaplan-Meier plot for the samples with mutations in S vs samples with no mutation in S . Nodes in network are color coded according to their contribution to the normalized log-rank statistic of the network (dark color = higher contribution). White nodes have no mutations. The p -value from the log-rank test is shown. (b) Subnetwork S associated with survival in OV, and Kaplan-Meier plot for the samples with mutations in S vs samples with no mutation in S . Color-coding as in (a).

the same cancer type as well as novel genes, while no gene is significantly associated with survival when considered in isolation.

There are many directions in which this work can be extended. First, we only considered single nucleotide variants and indels in our analysis; we plan to extend our method to consider more complex variants (e.g., copy number aberrations and differential methylation) in the analysis. Second, we believe that our algorithm and its analysis could be extended to the identification of subnetworks associated with clinical parameters other than survival time and to case-control studies, but substantial modifications to the algorithm and to its analysis will be required. Third, this work considers the log-rank statistic as a measure of association with survival; another popular test in survival analysis is the use of Cox’s regression model [21]. The two tests are identical in the case of two populations, therefore our algorithm identifies subnetworks with high score w.r.t. Cox’s regression model as well. However, Cox’s regression model allows for the correction for covariates (e.g., gender, age, etc.) in the analysis of survival data. A similar approach could be obtained by stratifying the patients in the log-rank test, but how to efficiently identify subnetworks, and in general combinations of genomic features, associated with survival while correcting for covariates remains a challenging open problem.

Finally, in this work we have restricted NoMAS to look for subnetworks of size at most 8, due to the fact that NoMAS is exponential in k , and that the running time for identifying larger subnetworks with low error probabilities therefore is expensive (albeit significantly faster than an exhaustive enumeration). However, we see that subnetworks of size k reported by NoMAS are very likely to contain genes that overlap with reported subnetworks of size $< k$. Thus, solving a smaller problem, which is cheaper to compute, might provide us with parts of the solution to a bigger problem, i.e. indicate vertices of the gene-gene interaction network that are of interest when searching for larger subnetworks. We can define a local search space $V' \subset V$ around such interesting *seed* vertices, for example as consisting of all the vertices reachable by at most t edges from any seed vertex, for some parameter t specifying the size/diversity of the search space. By adjusting t (as well as the number of seed vertices), this local search approach within a reduced vertex set will allow NoMAS to more efficiently find subnetworks of the current maximum size, but also to look for even larger subnetworks. The greedy algorithms are very fast and might therefore be good candidates for finding seed vertices. However, other methods for finding interesting areas of the gene-gene interaction network, such as using (or even combining) high scoring subnetworks identified by NoMAS for smaller problem sizes, can also be explored.

5 Acknowledgements

This work is supported, in part, by MIUR of Italy under project AMANDA and by NSF grant IIS-1247581. The results presented in this manuscript are in whole or part based upon data generated by the TCGA Research Network: <http://cancergenome.nih.gov/>. This paper was selected for oral presentation at RECOMB 2016 and an abstract is published in the conference proceedings

References

- [1] Noga Alon, Phuong Dao, Iman Hajirasouliha, Fereydoun Hormozdiari, and S Cenk Sahinalp. Biomolecular network motif counting and discovery by color coding. *Bioinformatics*, 24(13):i241–9, Jul 2008.
- [2] Noga Alon, Raphy Yuster, and Uri Zwick. Color-coding: a new method for finding simple paths, cycles and other small subgraphs within large graphs. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 326–335. ACM, 1994.
- [3] Cameron W Brennan, Roel G W Verhaak, Aaron McKenna, Benito Campos, Houtan Noushmehr, Sofie R Salama, Siyuan Zheng, Debyani Chakravarty, J Zachary Sanborn, Samuel H Berman, Rameen Beroukhi, Brady Bernard, Chang-Jiun Wu, Giannicola Genovese, Ilya Shmulevich, Jill Barnholtz-Sloan, Lihua Zou, Rahulsimham Vegesna, Sachet A Shukla, Giovanni Ciriello, W K Yung, Wei Zhang, Carrie Sougnez, Tom Mikkelsen, Kenneth Aldape, Darell D Bigner, Erwin G Van Meir, Michael Prados, Andrew Sloan, Keith L Black, Jennifer Eschbacher, Gaetano Finocchiaro, William Friedman, David W Andrews, Abhijit Guha, Mary Iacocca, Brian P O’Neill, Greg Foltz, Jerome Myers, Daniel J Weisenberger, Robert Penny, Raju Kucherlapati, Charles M Perou, D Neil Hayes, Richard Gibbs, Marco Marra, Gordon B Mills, Eric Lander, Paul Spellman, Richard Wilson, Chris Sander, John Weinstein, Matthew Meyerson, Stacey Gabriel, Peter W Laird, David Haussler, Gad Getz, Lynda Chin, and TCGA Research Network. The somatic genomic landscape of glioblastoma. *Cell*, 155(2):462–77, Oct 2013.
- [4] Sharon Bruckner, Falk Hüffner, Richard M Karp, Ron Shamir, and Roded Sharan. Topology-free querying of protein interaction networks. *J Comput Biol*, 17(3):237–52, Mar 2010.
- [5] Cancer Genome Atlas Network. Comprehensive genomic characterization of head and neck squamous cell carcinomas. *Nature*, 517(7536):576–82, Jan 2015.
- [6] Cancer Genome Atlas Research Network. Integrated genomic analyses of ovarian carcinoma. *Nature*, 474(7353):609–15, Jun 2011.
- [7] Cancer Genome Atlas Research Network. Comprehensive molecular characterization of clear cell renal cell carcinoma. *Nature*, 499(7456):43–9, Jul 2013.
- [8] Cancer Genome Atlas Research Network. Integrated genomic characterization of papillary thyroid carcinoma. *Cell*, 159(3):676–90, Oct 2014.
- [9] Daniel V T Catenacci, Gustavo Cervantes, Soheil Yala, Erik A Nelson, Essam El-Hashani, Rajani Kanteti, Mohamed El Dinali, Rifat Hasina, Johannes Brägelmann, Tanguy Seiwert, Michele Sanicola, Les Henderson, Tatyana A Grushko, Olufunmilayo Olopade, Theodore Karrison, Yung-Jue Bang, Woo Ho Kim, Maria Tretiakova, Everett Vokes, David A Frank, Hedy L Kindler, Heather Huet, and Ravi Salgia. Ron (mst1r) is a novel prognostic marker and therapeutic target for gastroesophageal adenocarcinoma. *Cancer Biol Ther*, 12(1):9–46, Jul 2011.
- [10] Salim A Chowdhury, Rod K Nibbe, Mark R Chance, and Mehmet Koyutürk. Subnetwork state functions define dysregulated subnetworks in cancer. *Journal of Computational Biology*, 18(3):263–281, 2011.
- [11] Han-Yu Chuang, Eunjung Lee, Yu-Tsueng Liu, Doheon Lee, and Trey Ideker. Network-based classification of breast cancer metastasis. *Mol Syst Biol*, 3:140, 2007.

- [12] Phuong Dao, Kendric Wang, Colin Collins, Martin Ester, Anna Lapuk, and S Cenk Sahinalp. Optimally discriminative subnetwork markers predict response to chemotherapy. *Bioinformatics*, 27(13):i205–i213, 2011.
- [13] Jishnu Das and Haiyuan Yu. Hint: High-quality protein interactomes and their applications in understanding human disease. *BMC Syst Biol*, 6:92, 2012.
- [14] Nathan D Dees, Qunyan Zhang, Cyriac Kandoth, Michael C Wendl, William Schierding, Daniel C Koboldt, Thomas B Mooney, Matthew B Callaway, David Dooling, Elaine R Mardis, Richard K Wilson, and Li Ding. Music: identifying mutational significance in cancer genomes. *Genome Res*, 22(8):1589–98, Aug 2012.
- [15] Alessandra Fiorio Pla and Dimitra Gkika. Emerging role of trp channels in cell migration: from tumor vascularization to metastasis. *Front Physiol*, 4:311, 2013.
- [16] Levi A Garraway and Eric S Lander. Lessons from the cancer genome. *Cell*, 153(1):17–37, Mar 2013.
- [17] Deliang Guo, Erica Hlavin Bell, and Arnab Chakravarti. Lipid metabolism emerges as a promising target for malignant glioma therapy. *CNS Oncol*, 2(3):289–99, May 2013.
- [18] Marc-Eric Halatsch, Sarah Löw, Kay Mursch, Thomas Hielscher, Ursula Schmidt, Andreas Unterberg, Vassilios I Vougioukas, and Friedrich Feuerhake. Candidate genes for sensitivity and resistance of human glioblastoma multiforme cell lines to erlotinib. laboratory investigation. *J Neurosurg*, 111(2):211–8, Aug 2009.
- [19] Matan Hofree, John P Shen, Hannah Carter, Andrew Gross, and Trey Ideker. Network-based stratification of tumor mutations. *Nat Methods*, 10(11):1108–15, Nov 2013.
- [20] Fereydoun Hormozdiari, Osnat Penn, Elhanan Borenstein, and Evan E Eichler. The discovery of integrated gene networks for autism and related disorders. *Genome Res*, 25(1):142–54, Jan 2015.
- [21] John Kalbfleisch and Ross Prentice. *The Statistical Analysis of Failure Time Data (Wiley Series in Probability and Statistics)*. Wiley-Interscience, 2 edition, 2002.
- [22] Cyriac Kandoth, Michael D McLellan, Fabio Vandin, Kai Ye, Beifang Niu, Charles Lu, Mingchao Xie, Qunyan Zhang, Joshua F McMichael, Matthew A Wyczalkowski, Mark D M Leiserson, Christopher A Miller, John S Welch, Matthew J Walter, Michael C Wendl, Timothy J Ley, Richard K Wilson, Benjamin J Raphael, and Li Ding. Mutational landscape and significance across 12 major cancer types. *Nature*, 502(7471):333–9, Oct 2013.
- [23] Yoo-Ah Kim, Dong-Yeon Cho, Phuong Dao, and Teresa M Przytycka. Memcover: integrated analysis of mutual exclusivity and functional network reveals dysregulated pathways across multiple cancer types. *Bioinformatics*, 31(12):i284–i292, 2015.
- [24] Michael S Lawrence, Petar Stojanov, Paz Polak, Gregory V Kryukov, Kristian Cibulskis, Andrey Sivachenko, Scott L Carter, Chip Stewart, Craig H Mermel, Steven A Roberts, Adam Kiezun, Peter S Hammerman, Aaron McKenna, Yotam Drier, Lihua Zou, Alex H Ramos, Trevor J Pugh, Nicolas Stransky, Elena Helman, Jaegil Kim, Carrie Sougnez, Lauren Ambrogio, Elizabeth Nickerson, Erica Shefler, Maria L Cortés, Daniel Auclair, Gordon Saksena, Douglas Voet, Michael Noble, Daniel DiCara, Pei Lin, Lee Lichtenstein, David I Heiman, Timothy Fennell, Marcin Imielinski, Bryan Hernandez, Eran Hodis, Sylvan Baca, Austin M Dulak, Jens Lohr, Dan-Avi Landau, Catherine J Wu, Jorge Melendez-Zajgla, Alfredo Hidalgo-Miranda, Amnon Koren, Steven A McCarroll, Jaume Mora, Ryan S

- Lee, Brian Crompton, Robert Onofrio, Melissa Parkin, Wendy Winckler, Kristin Ardlie, Stacey B Gabriel, Charles W M Roberts, Jaclyn A Biegel, Kimberly Stegmaier, Adam J Bass, Levi A Garraway, Matthew Meyerson, Todd R Golub, Dmitry A Gordenin, Shamil Sunyaev, Eric S Lander, and Gad Getz. Mutational heterogeneity in cancer and the search for new cancer-associated genes. *Nature*, 499(7457):214–8, Jul 2013.
- [25] Mark D M Leiserson, Fabio Vandin, Hsin-Ta Wu, Jason R Dobson, Jonathan V Eldridge, Jacob L Thomas, Alexandra Papoutsaki, Younhun Kim, Beifang Niu, Michael McLellan, Michael S Lawrence, Abel Gonzalez-Perez, David Tamborero, Yuwei Cheng, Gregory A Ryslik, Nuria Lopez-Bigas, Gad Getz, Li Ding, and Benjamin J Raphael. Pan-cancer network analysis identifies combinations of rare somatic mutations across pathways and protein complexes. *Nat Genet*, 47(2):106–14, Feb 2015.
- [26] Mark D M Leiserson, Hsin-Ta Wu, Fabio Vandin, and Benjamin J Raphael. Comet: a statistical approach to identify combinations of mutually exclusive alterations in cancer. *Genome Biol*, 16:160, 2015.
- [27] Shuli Liang, Lijie He, Xiaodi Zhao, Yu Miao, Yong Gu, Changcun Guo, Zengfu Xue, Weijia Dou, Fengrong Hu, Kaichun Wu, Yongzhan Nie, and Daiming Fan. MicroRNA let-7f inhibits tumor invasion and metastasis by targeting myh9 in human gastric cancer. *PLoS One*, 6(4):e18409, 2011.
- [28] N. Mantel. Evaluation of survival data and two new rank order statistics arising in its consideration. *Cancer Chemother Rep*, 50:163–170, 1966.
- [29] Sean Maxwell, Mark R Chance, and Mehmet Koyutürk. Efficiently enumerating all connected induced subgraphs of a large molecular network. In *Algorithms for Computational Biology*, pages 171–182. Springer, 2014.
- [30] D Williams Parsons, Siân Jones, Xiaosong Zhang, Jimmy Cheng-Ho Lin, Rebecca J Leary, Philipp Angenendt, Parminder Mankoo, Hannah Carter, I-Mei Siu, Gary L Gallia, Alessandro Olivi, Roger McLendon, B Ahmed Rasheed, Stephen Keir, Tatiana Nikolskaya, Yuri Nikolsky, Dana A Busam, Hanna Tekleab, Luis A Diaz, Jr, James Hartigan, Doug R Smith, Robert L Strausberg, Suely Kazue Nagahashi Marie, Sueli Mieko Oba Shinjo, Hai Yan, Gregory J Riggins, Darell D Bigner, Rachel Karchin, Nick Papadopoulos, Giovanni Parmigiani, Bert Vogelstein, Victor E Velculescu, and Kenneth W Kinzler. An integrated genomic analysis of human glioblastoma multiforme. *Science*, 321(5897):1807–12, Sep 2008.
- [31] Vishal N Patel, Giridharan Gokulrangan, Salim A Chowdhury, Yanwen Chen, Andrew E Sloan, Mehmet Koyutürk, Jill Barnholtz-Sloan, and Mark R Chance. Network signatures of survival in glioblastoma multiforme. *PLoS Comput Biol*, 9(9):e1003237, 2013.
- [32] R. Peto and J. Peto. Asymptotically efficient rank invariant test procedures. *J Roy Stat Soc Ser A*, 135:185+, 1972.
- [33] Benjamin J Raphael, Jason R Dobson, Layla Oesper, and Fabio Vandin. Identifying driver mutations in sequenced cancer genomes: computational approaches to enable precision medicine. *Genome Med*, 6(1):5, 2014.
- [34] Jüri Reimand and Gary D Bader. Systematic analysis of somatic mutations in phosphorylation signaling predicts novel cancer drivers. *Mol Syst Biol*, 9:637, 2013.
- [35] Raunak Shrestha, Ermin Hodzic, Jake Yeung, Kendric Wang, Thomas Sauerwald, Phuong Dao, Shawn Anderson, Himisha Beltran, Mark A Rubin, Colin C Collins, et al. Hit’ndrive: Multi-driver gene

- prioritization based on hitting time. In *Research in Computational Molecular Biology*, pages 293–306. Springer, 2014.
- [36] Padhraic Smyth and Rodney M Goodman. An information theoretic approach to rule induction from databases. *Knowledge and Data Engineering, IEEE Transactions on*, 4(4):301–316, 1992.
 - [37] David Tamborero, Abel Gonzalez-Perez, Christian Perez-Llamas, Jordi Deu-Pons, Cyriac Kandoth, Jüri Reimand, Michael S Lawrence, Gad Getz, Gary D Bader, Li Ding, and Nuria Lopez-Bigas. Comprehensive identification of mutational cancer driver genes across 12 tumor types. *Sci Rep*, 3:2650, 2013.
 - [38] Fabio Vandin, Patrick Clay, Eli Upfal, and Benjamin J Raphael. Discovery of mutated subnetworks associated with clinical data in cancer. *Pac Symp Biocomput*, pages 55–66, 2012.
 - [39] Fabio Vandin, Alexandra Papoutsaki, Benjamin J Raphael, and Eli Upfal. Accurate computation of survival statistics in genome-wide studies. *PLoS Comput Biol*, 11(5):e1004071, May 2015.
 - [40] Fabio Vandin, Eli Upfal, and Benjamin J Raphael. De novo discovery of mutated driver pathways in cancer. *Genome Res*, 22(2):375–85, Feb 2012.
 - [41] Bert Vogelstein, Nickolas Papadopoulos, Victor E Velculescu, Shibin Zhou, Luis A Diaz, Jr, and Kenneth W Kinzler. Cancer genome landscapes. *Science*, 339(6127):1546–58, Mar 2013.
 - [42] Guanming Wu and Lincoln Stein. A network module-based method for identifying cancer prognostic signatures. *Genome Biol*, 13(12):R112, 2012.
 - [43] Da Yang, Sofia Khan, Yan Sun, Kenneth Hess, Ilya Shmulevich, Anil K Sood, and Wei Zhang. Association of brca1 and brca2 mutations with survival, chemotherapy sensitivity, and gene mutator phenotype in patients with ovarian cancer. *JAMA*, 306(14):1557–65, Oct 2011.
 - [44] Haiyuan Yu, Leah Tardivo, Stanley Tam, Evan Weiner, Fana Gebreab, Changyu Fan, Nenad Svrzikapa, Tomoko Hirozane-Kishikawa, Edward Rietman, Xinpeng Yang, Julie Sahalie, Kourosh Salehi-Ashtiani, Tong Hao, Michael E Cusick, David E Hill, Frederick P Roth, Pascal Braun, and Marc Vidal. Next-generation sequencing to generate interactome datasets. *Nat Methods*, 8(6):478–80, Jun 2011.

Appendix

A Proof Sketches

Note that in Equation 1, the j -th term in the sum depends on the values of x_1, \dots, x_{j-1} . Our results use the fact that $V(\mathbf{x})$ can be computed as a sum of weights, one for each sample, that are independent of the value of the entries of \mathbf{x} .

Proposition 1. For $i = 1, 2, \dots, m$, let $w_i = c_i - \sum_{j=1}^i \frac{c_j}{m-j+1}$. Then

$$\sum_{j=1}^m c_j \left(x_j - \frac{m_1 - \sum_{i=1}^{j-1} x_i}{m-j+1} \right) = \sum_{j=1}^m w_j x_j.$$

Proof (Sketch).

$$\begin{aligned} \sum_{j=1}^m w_j x_j &= \sum_{j=1}^m \left(c_j - \sum_{i=1}^j \frac{c_i}{m-i+1} \right) x_j \\ &= \left(\sum_{j=1}^m c_j x_j \right) - \left(\sum_{j=1}^m x_j \sum_{i=1}^j \frac{c_i}{m-i+1} \right) \\ &= \left(\sum_{j=1}^m c_j x_j \right) - \left(\sum_{j:c_j \neq 0} \frac{\sum_{i=j}^m x_i}{m-j+1} \right) \\ &= \left(\sum_{j=1}^m c_j x_j \right) - \left(\sum_{j=1}^m c_j \frac{m_1 - \sum_{i=1}^{j-1} x_i}{m-j+1} \right) \\ &= \sum_{j=1}^m c_j \left(x_j - \frac{m_1 - \sum_{i=1}^{j-1} x_i}{m-j+1} \right) = V(\mathbf{x}). \end{aligned}$$

□

Theorem 1. The max k -set log-rank problem is NP-hard.

Proof (Sketch). The reduction is from the minimum set cover problem. In particular, we will show that if we can find a set \mathcal{S} with $|\mathcal{S}|$ maximizing $w'(\mathcal{S})$ in polynomial time, then we can test (in polynomial time) if there is a set cover of cardinality k . This implies that one could find the size of the minimum set cover in polynomial time, that is an NP-hard problem.

In the minimum set cover problem, one is given elements e_1, \dots, e_n , where each element $e_i, 1 \leq i \leq n$ is a subset of a universe set \mathcal{U} , with $|\mathcal{U}| = m$. The goal is to find the minimum cardinality subset $\mathcal{C} \subset \{e_1, \dots, e_n\}$ such that $\cup_{e \in \mathcal{C}} e = \mathcal{U}$.

Given an instance of the minimum set cover problem, we build an instance of the max k -set log-rank as follows. For each element $e_i, 1 \leq i \leq n$, we have a gene g_i , with $\mathcal{G} = \{g_1, \dots, g_n\}$. The set \mathcal{P} of patients has cardinality $4|\mathcal{U}|$. \mathcal{P} is partitioned into two sets \mathcal{P}_1 and \mathcal{P}_2 , with $\mathcal{P}_1 \cap \mathcal{P}_2 = \emptyset$ and $\mathcal{P} = \mathcal{P}_1 \cup \mathcal{P}_2$. Moreover we have $|\mathcal{P}_1| = \mathcal{U}$ and $|\mathcal{P}_2| = 3\mathcal{U}$, and the survival time of all patients in \mathcal{P}_1 is lower than the survival time of all patients in \mathcal{P}_2 . In addition, no patient of \mathcal{P}_1 is censored, while all patients in \mathcal{P}_2 are censored. The patients of \mathcal{P}_1 correspond to the elements of \mathcal{U} , and gene g_i is mutated in patients $e_i \subset \mathcal{U}$.

We now show that there is a minimum set cover of cardinality k if and only if $\max_{\mathcal{S} \subset \mathcal{G}, |\mathcal{S}|=k} w'(\mathcal{S}) = \frac{m}{4} - \sum_{j=1}^{m/4} \frac{1}{m-j+1}$. In particular, we will show that the maximum log-rank statistic is obtained when $x_i = 1$

for all $1 \leq i \leq \frac{m}{4}$ and $x_i = 0$ for all $\frac{m}{4} < i \leq n$, that can be achieved if and only if there is a set cover of cardinality k . (Note that m is divisible by 4 by construction.)

To prove the above, it is enough to show the following:

- i) $w_i > 0$ for $1 \leq i \leq \frac{m}{4}$;
- ii) for a fixed $m_1 \leq \frac{m}{4}$, the maximum weight is given by $\frac{\sum_{i=1}^{m_1} w_i}{\sqrt{m_1(m-m_1)}}$;
- iii) for all $1 \leq j \leq \frac{m}{4} - 1$: $\frac{\sum_{i=1}^j w_i}{\sqrt{j(m-j)}} \leq \frac{\sum_{i=1}^{j+1} w_i}{\sqrt{(j+1)(m-(j+1))}}$.

We first note that $w_i > w_{i+1}$ for $1 \leq i < \frac{m}{4}$.

$$w_i - w_{i+1} = 1 - \sum_{j=1}^i \frac{1}{m-j+1} - 1 + \sum_{j=1}^{i+1} \frac{1}{m-j+1} = \frac{1}{m-i} > 0.$$

To prove i) above it is then enough to prove that $w_{\frac{m}{4}} > 0$.

$$w_{\frac{m}{4}} = 1 - \sum_{j=1}^{\frac{m}{4}} \frac{1}{m-j+1} = 1 - \sum_{j=\frac{3m}{4}+1}^m \frac{1}{j} = 1 - H(m) + H\left(\frac{3m}{4}\right).$$

where $H(m)$ is the m -th harmonic number. Since $H(m) \leq \ln m + \gamma + \frac{1}{2m}$, with $\gamma \leq 0.58$ constant, for m large enough, and $H(m) \geq \ln m$, we have:

$$w_{\frac{m}{4}} \geq 1 - \ln m - \gamma - \frac{1}{2m} + \ln \frac{3m}{4} \geq 1 - \ln \frac{4}{3} - \gamma - \frac{1}{2m} > 0.1 - \frac{1}{m} > 0 \quad (2)$$

for m large enough.

ii) follows immediately from i) and from $w_i > w_{i+1}$ for $1 \leq i < \frac{m}{4}$ (since fixed m_1 , the denominator $\sqrt{m_1(m-m_1)}$ is fixed).

iii) can be proved by induction. □

Theorem 2. *The max connected k -set log-rank problem on graphs with at least one node of degree $O\left(n^{\frac{1}{c}}\right)$, where $c > 1$ is constant, is NP-hard.*

Proof (Sketch). Take an instance of set cover with n elements. We can “encode” it in the neighbours of a node of degree n in a graph with n^c vertices, where $c > 0$ is a constant, using the same scheme used for Theorem 1. All other vertices have no mutations. Note that the reduction is polynomial. □

Proposition 2. *For every $k \geq 3$ there is a family of instances of the max connected k -set log-rank problem and colorings for which OPT is not found by our algorithm even if it is colorful.*

Proof (Sketch). Let the number of samples be $n = 8(k-1)$. The censoring information \mathbf{c} is such that $c_i = 1$ for $1 \leq i \leq \frac{n}{4}$ and $c_j = 0$ for $\frac{n}{4} + 1 \leq j \leq n$. From Theorem 1 we get that all weights $w_i > 0$ for $1 \leq i \leq \frac{n}{4}$. Let \mathcal{I} be a tree with one internal vertex v_0 and $k+1$ leaf vertices $\{v_1, v_2, \dots, v_{k-1}, \bar{v}_1, \bar{v}_2\}$. Consider a coloring \mathcal{C} in which $\mathcal{C}(v_i)$ are distinct for $0 \leq i \leq k-1$ and $\mathcal{C}(v_j) = \mathcal{C}(\bar{v}_j)$ for $1 \leq j \leq 2$. Let $\sigma(v)$ be the set of weights for vertex v , i.e containing a weight for each sample mutated in the gene associated with v . Assign the weights such that $\sigma(v_0) = \emptyset$, $\sigma(v_i) = \{w_i, w_{k-1+i}\}$ and $\sigma(\bar{v}_i) = \{w_1, w_2\}$. Note that for any $k \geq 3$ the optimal connected subnetwork $OPT = \mathcal{S} = \{v_0, v_1, \dots, v_{k-1}\}$ since $\sigma(\mathcal{S}) = \{w_1, w_2, \dots, w_{n/4}\}$. By construction OPT is colorful.

The idea of the construction is to have two *bad* colors. A color c is bad if it is assigned to two vertices. The vertex in OPT with color c is a *good* vertex, while the vertex with color c not in OPT is a *bad* vertex. In our construction v_1 and v_2 are good vertices and \bar{v}_1 and \bar{v}_2 are bad vertices. Recall that our algorithm combines two subnetworks that are connected by an edge, thus every subnetwork of size ℓ must be a combination of a leaf v_i and some subnetwork $W(T, v_0)$ of size $\ell - 1$. To generate OPT , at some point we will have that v_i is one of the good vertices while $W(T, v_0)$ contains the other good vertex. We will show that this cannot happen. In particular we argue that $W(T, v_0)$ cannot contain only one bad color and be a subset of OPT . Without loss of generality, assume v_1 is the vertex with a bad color in $W(T, v_0)$. Consider the time it is added to $W(T, v_0)$ by combination of some $W(Q, v_0) \setminus \{v_1, v_2\}$ and $W(\{\mathcal{C}(v_1)\}, v_1)$. However, our algorithm will choose to combine with \bar{v}_1 in stead of v_1 because \bar{v}_1 yields the largest increase in the normalized log-rank statistic. To see this, note that v_1 and \bar{v}_1 both add two weights to $\sigma(W(Q, v_0))$ that are not already in $\sigma(W(Q, v_0))$. Both options therefore have the same number of mutations, and their normalized log-rank statistic can be compared by simply comparing their log-rank statistic. By construction $\sigma(\bar{v}_0)$ contains the two largest weights, hence it yields the larger log-rank statistic. \square

Theorem 3. *For any optimal colorful connected subnetwork \mathcal{S} of size $k \geq 3$ and any algorithm \mathcal{A} which obtains subnetworks with colorsets of cardinality i by combining 2 subnetworks with colorsets of cardinality $< i$, by adding 3 neighbors to \mathcal{S} we have that \mathcal{A} may not discover \mathcal{S} .*

sketch. Let the k vertices of OPT be deemed *good* vertices. For each of three of the vertices in OPT we add a *bad* copy, so that the good vertex v and the bad vertex \bar{v} have the same color and the same connectivity to the vertices in $OPT \setminus \{v\}$. By definition of \mathcal{A} , \mathcal{S} is found by combining two subnetworks of cardinality $< k$, and because there are three good vertices in OPT , one of these subnetworks of cardinality $< k$ will contain at least two good vertices. We show that an evil adversary can ensure that two subnetworks \mathcal{S}_1 and \mathcal{S}_2 , both being entries in W and each containing a good vertex, will never be combined by \mathcal{A} .

The combination of \mathcal{S}_1 and \mathcal{S}_2 will happen across a specific edge in the graph between one vertex $v_1 \in \mathcal{S}_1$ and one vertex $v_2 \in \mathcal{S}_2$. If v_2 is a good vertex then there will be another subnetwork $\bar{\mathcal{S}}_2$ in W with the same colorset as \mathcal{S}_2 , namely in the column corresponding to the bad vertex \bar{v}_2 , and since the connectivities of v_2 and \bar{v}_2 to OPT are the same, \mathcal{A} must select one of them. Due to the fact that $|\mathcal{S}_1 \cup \mathcal{S}_2| < k$ the adversary will be able to plant mutations so that $\bar{\mathcal{S}}_2$ is chosen over \mathcal{S}_2 . If v_2 is neither a good nor a bad vertex the same argument can be made to show that the adversary can ensure that \mathcal{S}_2 will not contain any good vertices. \square

The following is a result that we need to prove the performance of NoMAS under the Planted Subnetwork Model.

Proposition 3. *For every censoring vector c : $\sum_{i=1}^m w_i = 0$.*

Proof (Sketch). When $c_i = 1$ for all $1 \leq i \leq m$, then we have

$$\begin{aligned}
\sum_{i=1}^m w_i &= \sum_{i=1}^m \left(c_i - \sum_{j=1}^i \frac{c_j}{m-j+1} \right) \\
&= \sum_{i=1}^m \left(1 - \sum_{j=1}^i \frac{1}{m-j+1} \right) \\
&= m - \sum_{i=1}^m \sum_{j=1}^i \frac{1}{m-j+1} \\
&= m - \sum_{i=1}^m i \frac{1}{i} \\
&= m - m \\
&= 0.
\end{aligned}$$

When one c_i is switched to the value 0, we have that the weight changes by a factor:

$$-1 + \sum_{j=i}^m \frac{1}{m-i+1} = 0 \quad (3)$$

where the -1 is subtracted to w_i , while the value $\frac{1}{m-i+1}$ is summed (i.e., not subtracted) to all terms w_j with $j \geq i$. Therefore, any change to the censoring vector leaves $\sum_{i=1}^m w_i = 0$. \square

Using the above, we can prove the following.

Theorem 4. Let M be a mutation matrix corresponding to m samples from the Planted Subnetwork Model. If $m \in \Omega(k^4(k+\varepsilon)\ln n)$ for a given constant $\varepsilon > 0$ and $O(\ln(1/\delta)e^k)$ color-coding iterations are performed, then our algorithm identifies the optimal solution \mathcal{D} to the max connected k -set log-rank with probability $\geq 1 - \frac{1}{n^\varepsilon} - \delta$.

Proof (Sketch). Assume that \mathcal{D} is colorful. We prove that if NoMAS has build a subnetwork (with $1 \leq i < k$ vertices) consisting of vertices of \mathcal{D} only, then if $m \in \Omega(k^2(k+\varepsilon)\ln n)$, NoMAS will expand such solution by only using vertices in \mathcal{D} . Since NoMAS starts to build solutions from each vertex in \mathcal{D} , this proves that NoMAS identifies the optimal solution. We show this by proving that any set $\mathcal{C} \subset \mathcal{G} \setminus \mathcal{D}$, when added to any subset $\mathcal{S} \subset \mathcal{D}$, does not provide an improvement in the score as just adding one of the genes in \mathcal{D} .

From the properties of the Planted Subnetwork Model (PSM), we have that if \mathcal{S} is a subset of \mathcal{D} , then $w(\mathcal{S}) \geq \frac{c'm}{k}$, where c' is a constant > 0 . For a set $\mathcal{C} \subset \mathcal{G} \setminus \mathcal{D}$, we can consider it as a “metagene” that is mutated with a certain probability q (constant) in each sample, where q depends on the genes in \mathcal{C} .

From Property 3, we have that $\mathbf{E}[w(\mathcal{S} \cup \mathcal{C}) - w(\mathcal{S})] = -qw(\mathcal{S}) \leq -q\frac{c'm}{k}$, since the sum of all weights w_i is 0 and \mathcal{C} adds weights from a set of weights that must sum to $-w(\mathcal{S})$. From the properties of PSM, for a gene $g \in \mathcal{D} \setminus \mathcal{S}$ we have $w(\mathcal{S} \cup \{g\}) - w(\mathcal{S}) \geq \frac{c''m}{k}$, with $c'' > 0$ constant. Note that $w(\mathcal{S} \cup \mathcal{C}) - w(\mathcal{S})$ is the sum of independent random variables, and each random variable can change the value of $w(\mathcal{S} \cup \mathcal{C}) - w(\mathcal{S})$ by a value $< m$. Moreover, the number of samples in which \mathcal{C} can have mutations while \mathcal{S} does not is at least $\frac{m}{k}$ and at most m . We can therefore use Hoeffding inequality to bound the probability that $w(\mathcal{S} \cup \mathcal{C}) > w(\mathcal{S} \cup \{g\})$ as follows:

$$\begin{aligned}
\Pr(w(\mathcal{S} \cup \mathcal{C}) > w(\mathcal{S} \cup \{g\})) &= \Pr(w(\mathcal{S} \cup \mathcal{C}) - w(\mathcal{S}) > w(\mathcal{S} \cup \{g\}) - w(\mathcal{S})) \\
&\leq e^{-d((\frac{m}{k})^2(\frac{m}{k})^2/m^3)} \\
&\leq \frac{1}{n^{k+\varepsilon}}
\end{aligned}$$

for an appropriate constant $d > 0$ and for $m \in \Omega(k^4(k + \varepsilon) \ln n)$. By union bound on all sets \mathcal{C} of cardinality $\leq k$, we have that $\Pr(w(\mathcal{S} \cup \mathcal{C}) > w(\mathcal{S} \cup \{g\})) \leq \frac{1}{n^{k+\varepsilon}} n^k = \frac{1}{n^\varepsilon}$. Therefore, when $m \in \Omega(k^4(k + \varepsilon) \ln n)$ and \mathcal{D} is colorful, then NoMAS finds \mathcal{D} with probability $\geq 1 - \frac{1}{n^\varepsilon}$. The probability that \mathcal{D} is not colorful in any of the $O(\ln(1/\delta)e^k)$ color-coding iterations is $\leq \delta$. Therefore, by union bound the probability that NoMAS does not identify \mathcal{D} when $m \in \Omega(k^4(k + \varepsilon) \ln n)$ is $\leq \delta + \frac{1}{n^\varepsilon}$, and the result follows. \square

B Modifications to NoMAS

We design two modifications of NoMAS that can solve some easy cases where NoMAS may not identify the highest scoring solution due to its subnetwork merging strategy:

- i) we merge a subnetwork $W(T, u)$ not only with subnetworks $W(R, v)$ where v is a neighbor of u , but with subnetworks $W(R, w)$ where w is a neighbor of *any* vertex in $W(T, u)$;
- ii) in $W(T, u)$, we store $\ell > 1$ different colorful subnetworks containing u and with colorset T , leading to $\leq \ell^2$ choices for combining two entries of W and a corresponding ℓ^2 increase in the time complexity of the algorithm.

We note that the time complexity required by modification i) above is still polynomial at most a factor $|V|^2/|E| \in \Omega(n)$ larger than that of NoMAS. We note that both modifications find the optimal solution in the problem instance of Proposition 1, while the second one will find the optimal solution in the problem instance of Theorem 3 if ℓ is large enough. The second modification was run using $\ell = 5$ in our experiments and storing in $W(T, u)$ the $\leq \ell$ highest scoring subnetworks in $\mathcal{S}'(T, u)$.

C Greedy algorithms

We considered three different greedy strategies for the max connected k -set log-rank problem. All three algorithms build solutions starting from each node $u \in G$ in iterations by adding nodes to the current solution \mathcal{S} , and differ in the way they enlarge the current subnetwork \mathcal{S} of size $1 \leq i < k$. The first, `Greedy1`, screens all vertices at distance 1 to \mathcal{S} and adds the one that results in the best subnetwork of size $i + 1$. The second, `GreedyK`, considers all vertices at distance $\leq k - i$ to \mathcal{S} , and enforces connectivity by greedily constructing a path from the selected vertex to a vertex in \mathcal{S} . The third, `GreedyDFS`, traverses shortest paths from \mathcal{S} to every vertex at distance $\leq k - i$ by a depth-first search. The vertices on some shortest path of length $j \leq k - i$ which improved \mathcal{S} the most are added to obtain a subnetwork of size $i + j$.

D Pseudo code for NoMAS

The pseudo code for NoMAS is divided into three algorithms. First, algorithm 1 highlights the overall color-coding scheme. Second, algorithm 2 describes how the dynamic programming table W is computed in order of increasing colorset group sizes. Finally, algorithm 3 details the process of computing the subnetwork at a specific entry in W . It is assumed that the undirected graph $G(V, E)$, the mutation matrix M and the survival information \mathbf{x}, \mathbf{c} are globally known. As a companion piece to algorithm 3, figure 4 visualizes the method used for combining two previously computed entries of W .

Algorithm 1: NoMAS(k, δ)

```

 $\mathcal{S} \leftarrow \emptyset$ 
for  $i \leftarrow 1$  to  $\ln(\frac{1}{\delta})e^k$  do
    Color the vertices of  $G$  with  $k$  colors uniformly at random
     $W \leftarrow \text{FILLTABLE}(k)$ 
     $\mathcal{S}' \leftarrow \arg \max_{\forall T \forall v : W(T, v) \in W} \{w(W(T, v))\}$ 
     $\mathcal{S} \leftarrow \arg \max \{w(\mathcal{S}), w(\mathcal{S}')\}$ 
return  $\mathcal{S}$ 

```

Algorithm 2: FILLTABLE(k)

```

 $W \leftarrow$  empty table with dimensions  $(2^k - 1) \times |V|$ 
for each vertex  $u \in V$  do
    for each color  $\alpha$  among the  $k$  colors do
        if the color of  $u$  is  $\alpha$  then
             $W(\{\alpha\}, u) \leftarrow \{u\}$ 
        else
             $W(\{\alpha\}, u) \leftarrow \emptyset$ 
    for  $i \leftarrow 2$  to  $k$  do
        /* The following may be distributed among  $N \leq |V|$  processors */
        for each vertex  $u \in V$  do
            for each colorset  $T$  of size  $i$  do
                 $W(T, u) \leftarrow \text{COMPUTEENTRY}(T, u)$ 
return  $W$ 

```

Algorithm 3: COMPUTEENTRY(T, u)

```

best  $\leftarrow \emptyset$ 
for each neighbor  $v$  of  $u$  do
    for each colorset  $Q$  s.t.  $Q \subset T$  and  $Q \neq \emptyset$  do
         $R \leftarrow T \setminus Q$ 
        if  $W(Q, u) \neq \emptyset$  and  $W(R, v) \neq \emptyset$  then
            candidate  $\leftarrow W(Q, u) \cup W(R, v)$ 
            best  $\leftarrow \arg \max \{w(\text{candidate}), w(\text{best})\}$ 
return best

```

Modifications The two proposed modifications to NoMAS differ from NoMAS in their method for computing an entry of W . Algorithm 4 describes modification i, while algorithm 5 details modification ii. Both algorithms should be seen as replacements for algorithm 3 of the unmodified version of NoMAS. Figure 5




Figure 4: Examples of several pairs of colorful connected subnetworks $W(Q, u)$ and $W(R, v)$ considered by NoMAS when computing the entry $W(T, u)$ for a colorset T of size 5. In each example a subnetwork containing u are combined with a subnetwork containing a neighbor v of u , in order to obtain a subnetwork with colorset $T = Q \cup R$ such that $Q \cap R = \emptyset$. The dotted edge is the one connecting the two subnetworks (the edge is always connected to u).

visualizes the combination strategy of algorithm 4 (note the difference from figure 4).

Algorithm 4: MODIFICATIONI(T, u)

```

best  $\leftarrow \emptyset$ 
for each colorset  $Q$  s.t.  $Q \subset T$  and  $Q \neq \emptyset$  do
   $R \leftarrow T \setminus Q$ 
  for each neighbor  $w$  of a vertex in  $W(Q, u)$  do
    candidate  $\leftarrow W(Q, u) \cup W(R, w)$ 
    best  $\leftarrow \arg \max\{w(\text{candidate}), w(\text{best})\}$ 
return best

```

Algorithm 5: MODIFICATIONII(T, u)

```

candidates  $\leftarrow \emptyset$ 
for each neighbor  $v$  of  $u$  do
  for each colorset  $Q$  s.t.  $Q \subset T$  and  $Q \neq \emptyset$  do
     $R \leftarrow T \setminus Q$ 
    for each subnetwork  $A \in W(Q, v)$  do
      for each subnetwork  $B \in W(R, v)$  do
        candidates  $\leftarrow \text{candidates} \cup \{A \cup B\}$ 
best  $\leftarrow$  the  $\ell$  distinct highest scoring subnetworks in candidates
return best

```






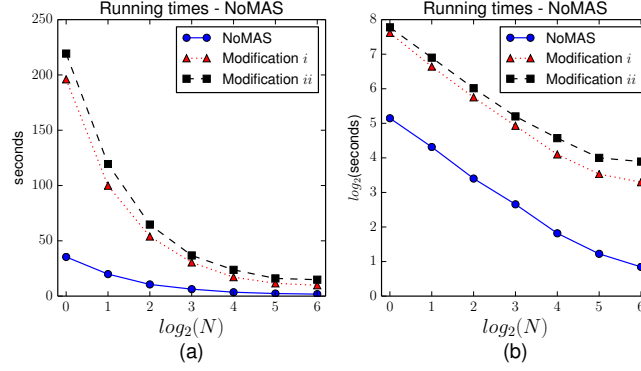
Figure 5: Examples of several pairs of colorful connected subnetworks $W(Q, u)$ and $W(R, w)$ considered by NoMAS with modification i when computing the entry $W(T, u)$ for a colorset T of size 5. In each example a subnetwork $W(Q, u)$ containing u are combined with a subnetwork containing a neighbor w of some vertex in $W(Q, u)$, in order to obtain a subnetwork with colorset $T = Q \cup R$ such that $Q \cap R = \emptyset$. The dotted edge is the one connecting the two subnetworks.




Supplementary Figure 1: p-values of the permutation test on the top 10 solutions identified by GreedyK for different values of k in both tail tests. The top 10 solutions on the permuted data are obtained using both GreedyK and NoMAS (with 32 color-coding iterations).



Supplementary Figure 2: Running time comparison between NoMAS and the exhaustive enumeration algorithm on three different cancer datasets. The running times of both algorithms are obtained using 40 processors. The running times for NoMAS account for 256 color-coding iterations and excludes the statistical assessment of the identified solutions.



Supplementary Figure 3: Running times of NoMAS and the two modifications considered for varying numbers of processors N . The running times are for a single iteration for $k = 8$ and are obtained on the OV cancer data (a) The running times in seconds. (b) The running times in seconds on a logarithmic scale.



Supplementary Figure 4: Running times of the three greedy algorithms and a single color-coding iteration of NoMAS for varying values of k and on three different cancer data. Each of the algorithms are run on a single processor. The top panels show the times measured when maximizing the score $w(S)$, while the bottom panels show the times for maximizing the score $-w(S)$.